

WEST[Help](#)[Logout](#)[Interrupt](#)[Main Menu](#)[Search Form](#)[Posting Counts](#)[Show S Numbers](#)[Edit S Numbers](#)[Preferences](#)[Cases](#)**Search Results -**

Terms	Documents
14 and 117	7

US Patents Full-Text Database
US Pre-Grant Publication Full-Text Database
JPO Abstracts Database
EPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Database:**Search:**[Refine Search](#)[Recall Text](#)[Clear](#)**Search History****DATE: Monday, May 13, 2002** [Printable Copy](#) [Create Case](#)

Set Name Query

side by side

DB=USPT,PGPB,JPAB,EPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>	<u>Set Name</u>
		result set	
<u>L22</u>	l4 and l17	7	<u>L22</u>
<u>L21</u>	l4 and l14	5	<u>L21</u>
<u>L20</u>	l4 and l13	5	<u>L20</u>
<u>L19</u>	((713/201)!.CCLS.))	976	<u>L19</u>
<u>L18</u>	((713/188)!.CCLS.))	57	<u>L18</u>
<u>L17</u>	((713/\$)!.CCLS.))	10403	<u>L17</u>
<u>L16</u>	((709/253)!.CCLS.))	136	<u>L16</u>
<u>L15</u>	((709/200)!.CCLS.))	451	<u>L15</u>
<u>L14</u>	((709/\$)!.CCLS.))	15667	<u>L14</u>
<u>L13</u>	((707/\$)!.CCLS.))	13901	<u>L13</u>
<u>L12</u>	((707/200)!.CCLS.))	857	<u>L12</u>
<u>L11</u>	((707/104.1)!.CCLS.))	1669	<u>L11</u>
<u>L10</u>	((707/100)!.CCLS.))	1013	<u>L10</u>
<u>L9</u>	((707/10)!.CCLS.))	1993	<u>L9</u>
<u>L8</u>	((707/1)!.CCLS.)	1451	<u>L8</u>
<u>L7</u>	L6 and backup and restor\$	4	<u>L7</u>
<u>L6</u>	second near transmit\$ near (mechanism or device)	419	<u>L6</u>
<u>L5</u>	L4 and first same communication and second same communication	12	<u>L5</u>
<u>L4</u>	L2 and backup and restor\$	24	<u>L4</u>
<u>L3</u>	L1 and "second communication"	8	<u>L3</u>
<u>L2</u>	L1 and "first computer"	79	<u>L2</u>
<u>L1</u>	"second computer" and backup and "data storage" and network	131	<u>L1</u>

END OF SEARCH HISTORY

WEST

L5: Entry 7 of 12

File: USPT

Dec 14, 1999

US-PAT-NO: 6003044

DOCUMENT-IDENTIFIER: US 6003044 A

TITLE: Method and apparatus for efficiently backing up files using multiple computer systems

DATE-ISSUED: December 14, 1999

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Pongracz; Gregory	Redwood City	CA		
Wertheimer; Steven	Kentfield	CA		
Bridge; William	Alameda	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Oracle Corporation	Redwood Shores	CA			02

APPL-NO: 8/ 962086 [PALM]

DATE FILED: October 31, 1997

PARENT-CASE:

RELATED APPLICATIONS The subject matter of this application is related to the subject matter of attorney docket number 1027, application Ser. No. 08/962,539 entitled, "METHOD AND APPARATUS FOR RESTORING A PORTION OF A DATABASE" filed on Oct. 31, 1997 by C. Gregory Doherty, Gregory Pongracz, William Bridge, Juan Loaiza and Mark Ramacher, attorney docket number 1028, application Ser. No. 08/962,087 entitled, "METHOD AND APPARATUS FOR IDENTIFYING FILES USED TO RESTORE A FILE" filed on Oct. 31, 1997 by Gregory Pongracz, Steven Wertheimer and William Bridge, attorney docket number 1036, application Ser. No. 08/961,747 entitled, "METHOD AND APPARATUS FOR PRESERVING NON-CURRENT INFORMATION THAT CAN BE OVERWRITTEN IN A COMPUTER FILE" filed on Oct. 31, 1997 by Gregory Pongracz and Tuomas Pystynen, attorney docket number 1038, application Ser. No. 08/961,741 entitled, "METHOD AND APPARATUS FOR ACCESSING A FILE THAT CAN BE CONCURRENTLY WRITTEN" filed on Oct. 31, 1997 by Tuomas Pystynen and Gregory Pongracz having the same assignee as this application and incorporated herein by reference in its entirety.

INT-CL: [6] G06 F 12/00

US-CL-ISSUED: 707/204; 711/162

US-CL-CURRENT: 707/204; 711/162

FIELD-OF-SEARCH: 707/204, 707/200, 707/208, 707/10, 395/826, 711/161, 711/162

PRIOR-ART-DISCLOSED:

U. S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>5515502</u>	May 1996	Wood	395/182.13
<input type="checkbox"/> <u>5721916</u>	February 1998	Pardikar	395/617
<input type="checkbox"/> <u>5790886</u>	August 1998	Allen	395/825
<input type="checkbox"/> <u>5799322</u>	August 1998	Mosher, Jr.	707/202
<input type="checkbox"/> <u>5819296</u>	October 1998	Anderson et al.	707/204
<input type="checkbox"/> <u>5832522</u>	November 1998	Blickenstaff et al.	707/204
<input type="checkbox"/> <u>5857193</u>	January 1999	Sutcliffe et al.	707/10
<input type="checkbox"/> <u>5857208</u>	January 1999	Ofek	707/204
<input type="checkbox"/> <u>5860122</u>	January 1999	Owada et al.	711/162

ART-UNIT: 277

PRIMARY-EXAMINER: Kulik, Paul V.

ATTY-AGENT-FIRM: Law Offices of Charles E. Gotlieb

ABSTRACT:

A system and method backs up computer files to backup drives connected to multiple computer systems. Each file in a backup set is allocated to one or more backup subsets for each of the multiple computer systems. The files can be allocated in an even number across each subset, allocated to evenly spread the number of bytes to each subset, or, using the capacity of each of the multiple computer systems, allocated so that each computer system can complete backing up the files allocated to it in approximately the same amount of time. The system can restrict the number of bytes continuously required from a single disk by one of the backup machines from exceeding a threshold limit. Each of the multiple computer systems is then directed to, and the computer systems do, back up files in one or more subsets, which may be allocated to that computer system.

17 Claims, 5 Drawing figures

WEST

End of Result Set

L5: Entry 12 of 12

File: USPT

Nov 23, 1971

US-PAT-NO: 3623014

DOCUMENT-IDENTIFIER: US 3623014 A

TITLE: COMPUTER COMMUNICATIONS SYSTEM

DATE-ISSUED: November 23, 1971

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Doelz; Melvin L.	Corona Del Mar	CA		
Martin; Robert F.	Orange	CA		
Morioka; Fred K.	Tustin	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Control Data Corporation	Santa Ana	CA			

APPL-NO: 4/ 852714 [PALM]

DATE FILED: August 25, 1969

INT-CL: [] G06f 15/16

US-CL-ISSUED: 340/172.5

US-CL-CURRENT: 714/6

FIELD-OF-SEARCH: 340/172.5, 235/157

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>3312954</u>	April 1967	Bible et al.	340/172.5
<input type="checkbox"/> <u>3374465</u>	March 1968	Richmond et al.	340/172.5
<input type="checkbox"/> <u>3377623</u>	April 1968	Reut et al.	340/172.5
<input type="checkbox"/> <u>3409877</u>	November 1968	Alterman et al.	340/172.5
<input type="checkbox"/> <u>3411139</u>	November 1968	Lynch et al.	340/172.5
<input type="checkbox"/> <u>3444528</u>	May 1969	Lovell et al.	340/172.5

ART-UNIT: 237

PRIMARY-EXAMINER: Shaw; Gareth D.

ASSISTANT-EXAMINER: Chapuran; Ronald F.

ATTY-AGENT-FIRM: Spensley, Horn and Lubitz

ABSTRACT:

A computer controlled, data communications system is provided for transmitting data between a plurality of external devices. The system comprises up to eight, general purpose, digital computers, each with an associated disc-file storage system. The computers and the disc-file storage units are organized such that communication between the computers is made via the disc-file storage associated with each computer, not directly between the computers themselves. Each group of external devices is coupled to a primary and secondary computer such that upon the failure of the primary computer, the secondary computer will process the data for its primary group of devices as well as for those external devices for which it is the secondary computer. Each disc-file storage unit is the primary storage for one computer and the copy storage for one additional computer. Upon the failure of a primary disc-file storage unit, a computer can operate with the copy storage unit.

7 Claims, 7 Drawing figures

WEST[Help](#)[Logout](#)[Interrupt](#)[Main Menu](#) [Search Form](#) [Posting Counts](#) [Show S Numbers](#) [Edit S Numbers](#) [Preferences](#) [Cases](#)**Search Results -**[Terms](#)[Documents](#)

L13 and port!

6

Database:

US Patents Full-Text Database
US Pre-Grant Publication Full-Text Database
JPO Abstracts Database
EPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search:[Refine Search](#)[Recall Text](#)[Clear](#)**Search History**DATE: **Monday, May 13, 2002** [Printable Copy](#) [Create Case](#)

Set Name Query
side by side

DB=USPT,PGPB,JPAB,EPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>	<u>Set Name</u>
		result set	
<u>L15</u>	L13 and port!	6	<u>L15</u>
<u>L14</u>	L13 and socket	6	<u>L14</u>
<u>L13</u>	L12 and database same application	24	<u>L13</u>
<u>L12</u>	l8 and network and restor\$	95	<u>L12</u>
<u>L11</u>	L10 and port!	11	<u>L11</u>
<u>L10</u>	L8 and socket	14	<u>L10</u>
<u>L9</u>	L8 and network near storage	1	<u>L9</u>
<u>L8</u>	L7 and backup and restore	107	<u>L8</u>
<u>L7</u>	first near computer and second near computer	5889	<u>L7</u>
<u>L6</u>	L5 and backup and restore	4	<u>L6</u>
<u>L5</u>	l4 and first near computer and second near computer	11	<u>L5</u>
<u>L4</u>	5206939.uref.	93	<u>L4</u>

DB=USPT; PLUR=YES; OP=OR

<u>L3</u>	4862411.pn.	1	<u>L3</u>
<u>L2</u>	4939598.pn.	1	<u>L2</u>

DB=USPT,PGPB,JPAB,EPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>L1</u>	5206939.pn.	3	<u>L1</u>
-----------	-------------	---	-----------

END OF SEARCH HISTORY

PALM INTRANET

Day : Monday
 Date: 5/13/2002
 Time: 07:14:35

Patent Number InformationApplication Number: **07/586796****Assignments**Filing Date: **09/24/1990**Examiner Number: **68908 / BOWLER, ALYSSA**Effective Date: **09/24/1990**Group Art Unit: **2312**Application Received: **09/24/1990**Class/Subclass: **395/400.000**Patent Number: **5206939**Lost Case: **NO**Issue Date: **04/27/1993**

Interference Number:

Date of Abandonment: **00/00/0000**

Unmatched Petition:

Attorney Docket Number: **EMC-103XX**L&R Code: Secrecy Code: **1**Status: **150 / PATENTED CASE**Third Level Review: **NO** Secrecy Order: **NO**Confirmation Number: **3295**Status Date: **04/15/1993**Title of Invention: **SYSTEM AND METHOD FOR DISK MAPPING AND DATA RETRIEVAL**

Bar Code #	Location	Location Date	Chrg to Location	Charge to Name	Emp. ID	Infra Loc
07586796	9210 RECORD ROOM 308-2733	05/07/2002			DWYATT	

Appln
Info

Search Another: Application# or Patent# PCT / / or PG PUBS # Attorney Docket # Bar Code #

(To Go BACK Use BACK Button on Your BROWSER Tool Bar)

Back to [PALM](#) | [ASSIGNMENT](#) | [OASIS](#) | Home page


PALM INTRANET

 Day : Monday
 Date: 5/13/2002
 Time: 07:12:56

Application Number Information

Application Number: 09/894422

Assignments

Filing Date: 06/28/2001

Examiner Number: 71586 / RINEHART, MARK

Effective Date: 06/28/2001

Group Art Unit: 2152

Application Received: 06/29/2001

Class/Subclass: 709/225.000

Patent Number:

Lost Case: NO

Issue Date: 00/00/0000

Interference Number:

Date of Abandonment: 00/00/0000

Unmatched Petition: NO

Attorney Docket Number: EMC-00-067

L&R Code: Secrecy Code:1

Status: 30 /DOCKETED NEW CASE - READY FOR EXAMINATION

Secrecy Order: NO

Confirmation Number: 6882

Status Date: 10/15/2001

**Title of Invention: SYSTEM AND METHOD FOR MANAGING REPLICATION SETS OF
DATA DISTRIBUTED OVER ONE OR MORE COMPUTER SYSTEMS**

Bar Code	Location	Location Date	Chrg to Loc	Charge to Name	Emp. ID	Infra Loc
09894422	2FC1 REPOSITORY CENTRAL FILES OVERFLOW FOR TC'S 1700, 2100, 2600 & 2800	03/20/2002		No Charge to Name	DWENDEMAGEGEH	FRAN/01/

[Appln](#) [Contents](#) [Petition Info](#) [Atty/Agent Info](#) [Continuity Data](#) [Foreign Data](#) [Inve](#)
[Info](#)
Search Another: Application# or Patent# PCT / / or PG PUBS # Attorney Docket # Bar Code #

(To Go BACK Use BACK Button on Your BROWSER Tool Bar)

Back to [PALM](#) | [ASSIGNMENT](#) | [OASIS](#) | Home page

**Application Number Information**

Application Number: 09/946007

Assignments

Filing Date: 09/04/2001

Examiner Number: 71586 / RINEHART, MARK

Effective Date: 09/04/2001

Group Art Unit: 2152

Application Received: 09/05/2001

Class/Subclass: 709/225.000

Patent Number:

Lost Case: NO

Issue Date: 00/00/0000

Interference Number:

Date of Abandonment: 00/00/0000

Unmatched Petition: NO

Attorney Docket Number: EMC-008PUS

L&R Code: Secrecy Code:1

Status: 30 /DOCKETED NEW CASE - READY FOR EXAMINATION

Secrecy Order: NO

Confirmation Number: 5827

Status Date: 03/24/2002

Title of Invention: **INFORMATION REPLICATION SYSTEM MOUNTING PARTIAL DATABASE REPLICATIONS**

Bar Code	Location	Location Date	Chrg to Loc	Charge to Name	Emp. ID	Infra Loc
09946007	2FC1 REPOSITORY CENTRAL FILES OVERFLOW FOR TC'S 1700, 2100, 2600 & 2800	03/28/2002		No Charge to Name	DWENDEMAGEGEH	FRAN/01/

Search Another: Application#

or Patent#

PCT / / or PG PUBS #

Attorney Docket #

Bar Code #

(To Go BACK Use BACK Button on Your BROWSER Tool Bar)

Back to [PALM](#) | [ASSIGNMENT](#) | [OASIS](#) | Home page

**Application Number Information**

Application Number: 09/946078

Assignments

Filing Date: 09/04/2001

Examiner Number: 71586 / RINEHART, MARK

Effective Date: 09/04/2001

Group Art Unit: 2152

Application Received: 09/05/2001

Class/Subclass: 709/225.000

Patent Number:

Lost Case: NO

Issue Date: 00/00/0000

Interference Number:

Date of Abandonment: 00/00/0000

Unmatched Petition: NO

Attorney Docket Number: EMC-007PUS

L&R Code: Secrecy Code:1

Status: 30 /DOCKETED NEW CASE - READY FOR EXAMINATION

Secrecy Order: NO

Confirmation Number: 4568

Status Date: 03/13/2002

Title of Invention: **INFORMATION REPLICATION SYSTEM HAVING ENHANCED ERROR DETECTION AND RECOVERY**

Bar Code	Location	Location Date	Chrg to Loc	Charge to Name	Emp. ID	Infra Loc
09946078	2FC1 REPOSITORY CENTRAL FILES OVERFLOW FOR TC'S 1700, 2100, 2600 & 2800	04/02/2002		No Charge to Name	DWENDEMAGEGEH	FRAN/01/

Appln
Info

Search Another: Application# or Patent# PCT / / or PG PUBS # Attorney Docket # Bar Code #

(To Go BACK Use BACK Button on Your BROWSER Tool Bar)

Back to PALM | ASSIGNMENT | OASIS | Home page

PALM INTRANET

Day : Monday
 Date: 5/13/2002
 Time: 07:16:27

Application Number InformationApplication Number: **09/000540****Assignments**Filing Date: **12/30/1997**Examiner Number: **76701 / PHAM, BRENDA**Effective Date: **12/30/1997**Group Art Unit: **2664**Application Received: **12/30/1997**Class/Subclass: **370/381.000**

Patent Number:

Waiting for Response Desc.
Mail Final Rej.Issue Date: **00/00/0000**Lost Case: **NO**Date of Abandonment: **00/00/0000**

Interference Number:

Attorney Docket Number: **EMC97029**Unmatched Petition: **NO**Status: **83 /ADVISORY ACTION MAILED**L&R Code: Secrecy Code:**1**Confirmation Number: **5453**Third Level Review: **NO**Secrecy Order: **NO**Title of Invention: **METHOD AND APPARATUS FOR TRANSFERS EMPLOYED IN DATA
STORAGE SYSTEM**Status Date: **02/12/2001**

Bar Code	Location	Location Date	Chrg to Loc	Charge to Name	Emp. ID	Infra Loc
09000540	26C1 TC 2600 CENTRAL FILES C1	11/07/2001		No Charge to Name	LOE CONV DR1995U	

Appln

Search Another: Application#

or Patent#

PCT / / or PG PUBS #

Attorney Docket #

Bar Code #

(To Go BACK Use BACK Button on Your BROWSER Tool Bar)

Back to [PALM](#) | [ASSIGNMENT](#) | [OASIS](#) | Home page

WEST

L1: Entry 1 of 3

File: USPT

Apr 27, 1993

US-PAT-NO: 5206939

DOCUMENT-IDENTIFIER: US 5206939 A

TITLE: System and method for disk mapping and data retrieval

DATE-ISSUED: April 27, 1993

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Yanai; Moshe	Framingham	MA		
Vishlitzky; Natan	Brookline	MA		
Alterescu; Bruno	Newton	MA		
Castel; Daniel	Farmingham	MA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
EMC Corporation	Hopkinton	MA			02

APPL-NO: 7/ 586796 [PALM]

DATE FILED: September 24, 1990

INT-CL: [5] G06F 13/00, G06F 12/06, G11B 5/09

US-CL-ISSUED: 395/400; 395/425, 364/DIG.1, 360/48

US-CL-CURRENT: 711/4; 360/48, 711/113, 711/206

FIELD-OF-SEARCH: 395/425, 395/595, 395/400, 360/48, 364/DIG.1, 364/252.3, 364/252.4, 364/255.2

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

	PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/>	<u>4262332</u>	April 1981	Bass et al.	364/200
<input type="checkbox"/>	<u>4310883</u>	January 1982	Clifton et al.	364/200
<input type="checkbox"/>	<u>4500954</u>	February 1985	Duke et al.	364/200
<input type="checkbox"/>	<u>4533955</u>	August 1985	Christian et al.	364/200
<input type="checkbox"/>	<u>4533996</u>	August 1985	Hartung et al.	364/200
<input type="checkbox"/>	<u>4862411</u>	August 1989	Dishon et al.	364/952
<input type="checkbox"/>	<u>4939598</u>	July 1990	Kulakowski et al.	360/48

ART-UNIT: 232

PRIMARY-EXAMINER: Bowler; Alyssa H.

ATTY-AGENT-FIRM: Weingarten, Schurigin, Gagnebin & Hayes

ABSTRACT:

An apparatus and method for disk mapping and data retrieval includes a data storage medium on which has been stored a plurality of data records. Each record includes at least a record identification portion, for uniquely identifying each record from among the plurality of data records. The apparatus builds a record locator table and high speed semiconductor memory which comprises the unique record identifiers for the records on the storage medium as well as a record locator index generated by the apparatus, which indicates the address of the data record on the storage medium. Data retrieval is facilitated by first searching the record locator table in high speed semiconductor memory for a requested data record. Utilizing the record locator index associated with the requested data record, the system directly accesses the requested data record on the storage medium thereby minimizing storage medium search time. Also disclosed is an apparatus and method for converting CKD formatted data records to FBA formatted disk drives and for building and compressing the "count" portion of the CKD data formatted record into a record locator table.

19 Claims, 7 Drawing figures
Exemplary Claim Number: 1
Number of Drawing Sheets: 6

BRIEF SUMMARY:

FIELD OF THE INVENTION

This invention relates to data storage on disk drives and more particularly, to a method and apparatus for retrieving data records stored on a storage medium utilizing a data record locator index stored in memory.

BACKGROUND OF THE INVENTION

Large disk storage systems like the 3380 and 3390 direct access storage devices (DASD) systems employed with many IBM mainframe computer systems are implemented utilizing many disk drives. These disk drives are specially made to implement a count, key, and data (CKD) record format on the disk drives. Disk drives utilizing the CKD format have a special "address mark" on each track signifying the beginning of a record on the track. After the address mark comes the three part record beginning with the "COUNT" which serves as the record ID and also indicates the lengths of both the optional key and the data portions of the record, followed by the optional "KEY" portion, which in turn is followed by the "DATA" portion of the record.

Although this format gives the system and user some flexibility and freedom in the usage of the disk drive, this flexibility forces the user to use more complicated computer programs for handling and searching data on the disk. Since the disk drive track has no physical position indicator, the disk drive controller has no idea of the data which is positioned under the read/write head at any given instant in time. Thus, before data can be read from or written to the disk drive, a search for the record must be performed by sequentially reading all the record ID's contained in the count field of all the records on a track until a match is found. In such a search, each record is sequentially searched until a matching ID is found. Even if cache memory is used, all the records to be searched must first be read into the cache before being searched. Since searching for the record takes much longer than actual data transfer, the disk storage system spends a tremendous amount of time searching for data which drastically reduces system performance.

Disk drives employing what is known as a Fixed Block Architecture (FBA) are widely available in small, high capacity packages. These drives, by virtue of, their architecture, tend to be of higher performance than drives employing a CKD format. Such FBA drives are available, for example, from Fujitsu as 5.25" drives with 1 gigabyte or greater capacity.

The distinct advantage of utilizing many small disk drives is the ability to form a disk array. Thus a large storage capacity can be provided in a reduced amount of space, and storage redundancy can be provided in a cost effective manner. A serious problem arises, however, when trying to do a "simple" conversion of data from CKD formatted disks to FBA disks. Two schemes for such a conversion have been considered

which do not provide an acceptable solution to the conversion problem. The first of such schemes involves placing every field i.e. Count, Key and Data, of the CKD formatted record into a separate block on the FBA disk drive. Although this scheme does not waste valuable disk space when CKD formatted records contain large amounts of data, the "Count" field which is very short (8 bytes) occupies an entire block which is typically at least 512 bytes. For example, a CKD formatted record containing 47K bytes of data could be converted to 95 blocks of FBA disk, 512 bytes in length. In such a conversion, one block would be used to store the count of the record while 94 blocks (47K bytes length of data divided 512 bytes of FBA disk block) would be used to store data, for a total of 95 blocks. However, search time for finding the desired record is still a problem since all the records must be sequentially searched.

For records having very short data lengths such as eight bytes, however, one full track, or 94 CKD formatted data records would need 188 blocks on the FBA disk: 94 blocks for the count portion of the records and 94 blocks for the data portion of the records, even though each data record may only occupy 8 bytes of a 512 byte FBA block. Such a scheme may thus waste nearly 50% of the disk space on an FBA disk drive.

The second scheme for converting data from CKD to FBA drives involves starting each CKD record in a separate block and then writing the complete record in sequential blocks. Utilizing such a scheme, the first FBA block will contain the "count" portion of the record as well as the optional key portion and the start of the data portion of the record. This scheme, however, produces serious system performance degradation when data must be written to the disk, since before writing data to the disk, the entire record must first be read into memory, modified, and subsequently written back to the disk drive. Such a loss in system performance is generally unacceptable.

SUMMARY OF THE INVENTION

This invention features an apparatus and method for retrieving one or more requested data records stored on a storage medium by searching for a data record identifier and associated data record locator index stored in high speed semiconductor memory. The apparatus receives one or more data records, each of the data records including at least a record identification portion and a data portion. The apparatus transfers and stores the data records to one or more data storage mediums. As the records are transferred to the data storage medium, the apparatus of the present invention generates a plurality of record locator indices, each of the record locator indices corresponding to one of the plurality of data records, for uniquely identifying the location of each of the data records stored on the storage medium.

The apparatus further includes high speed semiconductor memory for storing at least the plurality of record locator indices and the associated plurality of record identification portions. Upon receiving a request for one or more data records stored on the storage mediums, the apparatus of the present invention searches the high speed semiconductor memory utilizing the data record identification portion and locates the corresponding record locator index associated with the requested data record. The apparatus then directly retrieves the data record from the storage medium using the record locator index located during the search of semiconductor memory.

In the preferred embodiment, the data records are received in CKD format and stored on an FBA formatted disk drive. The record identification portions and associated record locator indices are combined to form one record locator table stored in one or more blocks of the FBA formatted disk drive and also copied in the high speed semiconductor memory.

A method for retrieving one or more requested data records stored on a storage medium is disclosed utilizing a data record locator index stored in memory and includes the steps of receiving a plurality of data records, each record including at least a record identification portion and the data portion, and transferring and storing the data records to one or more storage mediums. The method also includes generating a plurality of record locator indices, each of which are associated with one of the plurality of data records and uniquely identify the location of the each of the plurality of data records stored on the storage medium. Also included are the steps of storing at least a plurality of record locator indices and the associated plurality of record identification portions in memory. In response to a request for access to one or more of the plurality of data records, the method includes searching the memory, locating one or more data record identification portions and associated record locator indices corresponding to the one or more requested data

records, and directly retrieving from the storage medium the requested data records as directed by the record locator indicie.

In one embodiment, the method of the present invention includes transforming and encoding CKD formatted data records onto one or more FBA disk drives. Also in the preferred embodiment, the step of storing the data record to one or more storage medium, the step of storing further including the steps of transforming and encoding at least the record identification portion of each of the data records, generating a plurality of record locator indices, and combining the transformed and encoded record locator indices and record identification portions, for forming a record locator table stored in a high speed semiconductor memory.

DRAWING DESCRIPTION:

BRIEF DESCRIPTION OF THE DRAWINGS

These, and other features and advantages of the present invention are described below in the following detailed description and accompanying drawings, in which:

FIG. 1 is a block diagram of the system for disk mapping and data retrieval according to the present invention;

FIG. 2 is a schematic representation of a CKD formatted data record;

FIG. 3 is a schematic representation of several blocks from a fixed block disk drive in which has been inserted the CKD formatted data transformed according to the method of the present invention;

FIG. 4 is a detailed schematic representation of a portion of the data identification and locator table of FIG. 3;

FIG. 5 is a detailed schematic representation of a portion of the device identification and locator table of FIG. 4;

FIG. 6 is a flowchart of the method for transforming CKD formatted data into fixed block data including a method for preparing a record identification and locator table; and

FIG. 7 is a flowchart illustrating the method for compressing the length of the record identification and locator table according to the present invention.

DETAILED DESCRIPTION:

DETAILED DESCRIPTION OF THE INVENTION

In accordance with the present invention, the disk storage system 10, FIG. 1, for disk mapping and data retrieval includes one or more means for receiving write commands and data such as channel adapter boards 12a-12d. The channel adapter boards are adapted to receive disk read/write commands and data over a plurality of communication channels such as channels 14 from one or more host computers (not shown) over channels identified respectively as 1-8 in FIG. 1.

The channel adapter boards 12a-12d are connected to temporary or cache semiconductor memory storage unit 16 by means of bus 18. Bus 18 is also connected to one or more disk adapter boards 20 which read and write data to one or more disk drive units 22.

Each of the disk drive units 22 may include one or more disk drives, depending upon the user's requirements. Also included in the system is one or more uninterruptible power supply (UPS) 24.

In operation, one or more channel adapter boards 12a-12d receive write commands along with the accompanying data over one or more channels 14 from one or more host computers. In the preferred embodiment, the data is received in CKD format. In order to improve system performance, the disk storage system of the present invention does not wait for disk adapters 20 to locate and update the data on the appropriate disk drives but rather, the channel adapter boards store the data in CKD format in temporary semiconductor memory storage unit or cache 16.

In addition to storing the data that must be written to one or more disk drives 22, channel adapter boards 12a-12d store in the memory, an indication associated with each data record that must be written to disk, indicating to the disk adapters 20 that the associated data record stored in cache must be written to the disk drives.

A more detailed description of a disk storage system with write preservation utilizing a write pending indicator is described in U.S application Ser. No. 07/556,254 filed concurrently with the present application and incorporated herein by reference.

In one embodiment, the system and method for disk mapping and data retrieval includes mapping CKD formatted data onto fixed block disk drives. To facilitate understanding of the CKD to FBA data transformation a CKD record 30, FIG. 2 is shown and described below.

In order for a CKD disk drive to locate the first record on any given track, the disk drive read/write head must search the entire track until it encounters a position indicator called an address mark 32. Following a short gap 34 in the track, the first record 30 begins. The CKD record is divided into three fields or portions: the record identification portion, called the count, 36 followed by another gap, 37; the optional key portion 38; and the data portion 40.

The count portion of the data record uniquely identifies this record. The count is the portion of the record that a host system requesting access to a given record presents to the disk drive in order to enable the disk drive to search for and locate the record.

The count is comprised of 8 bytes of information. Bytes 0 and 1, as shown at 42, are used to designate the length of the data, while the third byte, 44, designates the length of the optional key field. The fourth byte 46, designates the record number on the track. The fifth and sixth bytes 49, and the seventh and eighth bytes 50-51, designate the head and cylinder numbers respectively, at which the record is located on the device. The second record 52, is located immediately following the end of the data portion 40 of the first record 30.

Thus, in accordance with the present invention, a portion of an FBA disk having CKD formatted data stored thereon is shown in FIG. 3 and includes record identification and locator table 58 including the 5 blocks labeled 60a-60e, for the tracks of the previous cylinder which is organized as described below.

The first data record 61 of the next CKD cylinder being emulated and located on the FBA drive begins in block, 62. If the data portion of the record is longer than the length of the second block 62, the data portion of the first record will be continued in the next block 64, and subsequent blocks as necessary to store the data portion of the first record. If the length of the first record 61 is equal to or smaller than block 62, any remaining unused portion of block 62 will not be used and instead, the data portion of record 2, 63, will begin in FBA disk drive block 64. This process will repeat itself until all of the blocks of a given cylinder being emulated have been copied. Thus, because the data portion of every record begins at the beginning of a block, (i.e. there is never more than one data record per block), once the system computes the address or block in which a requested data record resides, immediate access is possible with little or no disk drive search time. Utilizing the record identification and locator table of the present invention, the system is able to compute the number of fixed length blocks that must be read to retrieve all the data of a given record, as illustrated herein below. Each record identification and locator table is subsequently loaded into the system memory to facilitate and greatly reduce data record searching and data retrieval time.

Shown in greater detail in FIG. 4, is a data record and identification table 70 according to the present invention for one device such as one disk drive including multiple cylinders. Each cylinder being emulated includes such a table on the drive itself, as well as a corresponding copy in the semiconductor cache memory.

Each device record identification and locator table includes a header portion 72 followed by one or more cylinder portions 74-78. In turn, each cylinder portion is comprised of a plurality of track portions 80.

Device header portion 72 of the record locator and identification table includes such information as shown in Table 1 below, including the length of the header, line 1; the starting address of the device scratch memory address, line 2; and the length of the device header including the scratch area, line 3. The header also includes one or more bytes for device flags, line 4.

TABLE 1

DEVICE ID TABLE HEADER BUFFER AND FLAG OFFSETS

1. DV HEADER	SIZE	\$1000	LENGTH OF DV HEADER AT THE ID TABLE
2. DV SCRATCH	OFFSET		DV HEADER
			DV SCRATCH START ADDRESS
3. DV HEADER LENGTH	SIZE	\$1000	LENGTH OF DV HEADER AT THE ID TABLE
			INCLUDING THE SCRATCH AREA
DV HEADER BUFFERS OFFSETS			
4. DV FLAGS	OFFSET	0	DV TABLE FLAGS
5. DV STATISTICS	OFFSET	4	STATISTICS/RESERVED BYTES
6. DV READ TASK	OFFSET	\$40	READ TASK (ONLY ONE)
7. DV SENSE INFO	OFFSET	\$60	SENSE INFO FOR THIS DV
8. DV TABLE SELECT BUFFER	OFFSET	\$80	\$40 BYTES
			SELECTION BUFFER FOR THE DEVICE
9. RW COUNT BUFFER	OFFSET	\$C0	8 BYTES R/W COUNT COMMAND DATA BUFFER
10. DV WR PEND FLAGS	OFFSET	\$100	
		\$140	BYTES
			WR PENDING BIT PER CYLINDER
11. DV WR PEND GROUPS	OFFSET	\$240	
		\$20	BYTES
			WR PENDING BIT FOR 8 CYLINDERS
12. DV FLAGS SELECT BUFFER	OFFSET	\$280	
		\$40	BYTES
			SELECTION BUFFER FOR UPDATES
13. DV FMT CHANGED FLAGS	OFFSET	\$300	
		\$140	BYTES
			FMT CHANGED BIT FOR CYLINDER
14. DV FMT CHANGED GROUPS	OFFSET	\$440	
		\$20	BYTES
			FMT CHANGED BIT FOR 8 CYLINDERS
15. DV TEMP BLK	OFFSET	\$400	
		\$200	BYTES
			TEMP BLK FOR RECOVERY
16. TEMP CYL ID SLOT	OFFSET	\$600	
		\$A00	BYTES
			TEMP ID FOR RECOVERY

Such flags include a write pending flag which is set if one or more records on the device are temporarily stored in cache memory awaiting writing and storing to the disk drive, as well as an in-cache bit indicating that at least one record on the device is located in cache memory to speed up access to the record, line 4. Other bytes of the device header provide various informational, operational, or statistical data to the system.

For example, the write pending group flags shown at line 11 include one bit indicating a write pending on any one record on any of the 64 preselected consecutive cylinders comprising a cylinder group. Similarly, each cylinder has a write pending flag bit in the device header as shown at line 10. The various write pending flags form a "pyramid" or hierarchy of write pending flags which the system may search when no access to records stored on disks are requested for handling write pending requests. Such a hierarchy structure allows the system to inquire level by level within the structure whether any records on the device; any records within a group of cylinders; any records on a given cylinder; any records on a track; and record by record, whether any write pending flags are set or whether any records are located in cache memory. Such information is useful when processing data such as writing data to disk after a power failure as more fully described in U.S. patent application No. 07/586,254 filed concurrently herewith and incorporated herein by reference. The "last track" header is a header of a fictitious or non-existent track and serves to indicate that the record locator table was itself modified and must also be written to disk.

A more detailed description of the cylinder header portions 82a-82c of the device record identification and locator table 70 of FIG. 4 is shown in Table 2 wherein any given cylinder header includes such information as the length of the cylinder header, line 1; cylinder write pending flags, line 4; the physical address of the cylinder line 6; and the CRC error check byte of the cylinder, line 7.

TABLE 2

CYLINDER ID TABLE HEADER BUFFERS AND FLAGS OFFSETS	
LENGTH OF CYL HEADER AT	
1. CYL.sub.-- HEADER.sub.-- LENGTH	SIZE \$AO ID TABLE
2. CYL.sub.-- FLAG	
OFFSET	
0	CYL FLAGS
3. CYL.sub.-- FLAG.sub.-- AUX	
OFFSET	
1	ADD ON TO THE ABOVE
4. CYL.sub.-- WR.sub.-- PEND.sub.-- FLAGS	
OFFSET	
2	WR PENDING BIT PER TRACK
5. CYL.sub.-- STATISTICS	
OFFSET	
4	STATISTICS/RESERVED BYTES
6. CYL.sub.-- PH.sub.-- ADD	
OFFSET	
16	PH ADD OF CYL
7. CYL.sub.-- SLOT.sub.-- CRC	
OFFSET	
23	CRC OF CYL

Each track entry in the record identification and locator table is shown in greater detail in FIG. 5 and comprises a track header portion 84 and a track body portion 86. The track header portion 84 includes information as shown in Table 3 below, including a track flag byte, line 1; record-count bytes, line 2; the track CRC check byte, line 3; track compress patterns line 4; and cache address pointer, line 5.

The body 86, FIG. 5 of the track portion of the record identification and locator table includes a plurality of record flags 83, (line 6, Table 3,) beginning at byte 20 and record modifiers 85 (line 7, Table 3) beginning at byte 159 and extending sequentially backward as necessary or until a collision with the record flags 83 occurs.

TABLE 3

ID TABLE TRACK HEADER AND BODY OFFSETS
--

```

1. TRACK.sub.-- FLAG
    OFFSET
        0  TRACK FLAGS
2. .RECORD.sub.-- COUNT
    OFFSET
        1  NUMBER OF RECORDS AT THIS TRACK
3. TRACK.sub.-- CRC
    OFFSET
        5  CRC BYTE FOR TRACK
4. COMPRESS.sub.-- PATTERNS
    OFFSET
        6  TRACK COMPRESS PATTERNS
5. CACHE.sub.-- TRACK.sub.-- POINTER
    OFFSET
        14 POINTER TO CACHE
6. RECORD.sub.-- FLAGS
    OFFSET  COMMON FLAG POINTER
7. TRACK.sub.-- TABLE.sub.-- BODY
    OFFSET  TRACK BODY
    (MODIFIER)
    TRACK COMMON FLAG BITS
8. DEFECTIVE    BIT  7  DEFECTIVE TRACK
9. ALT          BIT  6  ALTERNATE TRACK
10. EX.sub.-- TRACK.sub.-- TABLE
    BIT  5  EXTENDED TRACK TABLE SLOT
11. WRT.sub.-- PEND
    BIT  4  WRITE PENDING IN TRACK
12. DIAG.sub.-- CYL
    BIT  3  DIAGNOSTICS CYL ('CE', 'SA')
13. NOT.sub.-- USED
    BIT  2  NOT USED
14. INVALID.sub.-- ID
    BIT  1  ID SLOT DEFECTIVE AND INVALID
15. IN.sub.-- CACHE
    BIT  0  TRACK IN CACHE FLAG
    RECORD FLAGS BITS
16. COMPRESS.sub.-- CODE
    BITS 0-3
        COMPRESS ALGO' FOR THIS
        RECORD
17. KEY.sub.-- IN.sub.-- CACHE
    BIT  4  KEY FIELD IN CACHE
18. DATA.sub.-- IN.sub.-- CACHE
    BIT  5  DATA FIELD IN CACHE
19. KEY.sub.-- W.sub.-- PEND
    BIT  6  KEY FIELD WRITE PENDING
20. DATA.sub.-- W.sub.-- PEND
    BIT  7  DATA FIELD WRITE PENDING

```

The track flags shown on line 1 in table 3 are described in detail on lines 8-15 and includes such bits indicating a defective track bit, line 8; a write pending bit, line 11; and a track in cache bit, line 15. Similarly, the record flag bits of line 6 are shown in greater detail in lines 16-20 including bits comprising the compression algorithm for this record, line 16; key and data fields in cache, lines 17 and 18; and key field and data field write pending bits, lines 19 and 20.

The channel adapters 12a-12d, FIG. 1, receive data and read/write commands from one or more hosts over its respective channels. The data records are provided by the host in CKD format and are stored in cache memory 16 in CKD format. All records stored in cache whether temporarily while awaiting writing to disk, or records which have been read from the disk to be stored in cache for quicker access, are stored in CKD format. When the record is to be written to the disk drive, one of disk adapters 20 reads the data from cache memory over bus 18 and converts the CKD formatted data to the format of the present invention including a record identifier and locator

table all of which can be stored in a plurality of fixed blocks before outputting the data over the disk adapters' SCSI interface to one or more of disk drives 22.

The present method for mapping CKD formatted data into fixed block disk drives is, in part, based on the recognition that under usual conditions, a sequence of CKD formatted records will include the "R" portion of the count identifying the record number from among a number of sequentially numbered data records of the same length. Further, the records are generally stored on the same device cylinder and accessed by the same device head. Additionally, the key length will generally be zero or some predetermined generally constant number. Thus, the method for disk mapping 100, FIG. 6 of the present invention includes establishing the profile of an expected record, step 110. In the preferred embodiment, the expected record is established with the count CCHH code as the physical cylinder and head identification, as well as the key length (K.sub.1)=0, data length (D.sub.1)=8 and the "R" byte of the count assigned as record number (n)=0. Further, the record flags are set to 00.

During step 112, the system employing the method of the present invention obtains the first CKD formatted record and compares the CKD record with the previously established expected record step 114. At step 116, a determination is made as to whether or not the CKD formatted record including the "count" and record flags match those of the expected data record. If the CKD formatted record and the expected record match, the method proceeds to step 118 wherein the body of the track portion of the record identification and locator table previously discussed in conjunction with FIG. 5 and table 3 is built. Since the CKD formatted record matched the previously established expected data record, the record flag is set to 00 and no entry is made in the record modifier portion of the track ID table. Subsequently, the "R" byte for the record number of the next expected data record is incremented by one, step 120, before returning to the step of obtaining the next CKD formatted data record at step 112.

If the results of the comparison at step 116 indicate that the CKD formatted record does not match the expected data record, the method proceeds to step 122 wherein a change format code (see Table 5) and record modifier, as required, are prepared. Next, the record format code and, if required, record modifier are inserted into the track identification table, step 124. If the track ID table is not full as determined at step 126, processing continues to step 128 wherein the current CKD count becomes the next expected count. Processing then returns to step 120 where the "R" byte is incremented by one before getting the next CKD record at step 112.

If, as indicated at step 126, the track identification table is full, meaning that the record flag portion of the ID table has collided with the record modifier portion of the ID table, the method of the instant invention will attempt to compress or shrink the body of the track ID table as shown in flowchart 130, FIG. 7. During the compression process of the instant invention, the system and method of the present invention attempt to define from one to eight data lengths which are repeated within this track. Such repeating data lengths are then classified as "patterns" and are thereafter referred to by a pattern "code" in the track header as shown on line 4 of Table 3, thus saving up to 2 bytes in the modifier portion of the track ID table for each repeated data length.

The method of the present invention first reads the ID table, step 132, searching for ID's with format code 03 for repeating values of data lengths, step 134. From those repeating values, the system and method of the present invention build a data length pattern table beginning with the data length that is most frequently repeated, and continuing on to find the seven most repeated data lengths and replaces the old 8 byte pattern with the new 8 byte pattern, step 136.

The method then proceeds to compare the data lengths of all the CKD records of the current track which have previously been read to determine whether or not any of the data lengths match the data patterns loaded in the pattern table, step 138. If any data lengths match the data patterns in the pattern table, the method proceeds to insert the data pattern code for the data length in a temporary ID table. Thus, the replaced record modifiers which previously contained the changed or modified data lengths are now unnecessary and eliminated, thus compressing or shrinking the size of the record identification and locator table and therefore allowing more room for the system to complete reading the CKD records for a given track. The system verifies at step 140, that the ID table was in fact compressed. If no ID table compression was achieved, the system reverts an error, step 142. If ID table compression was achieved, the method replaces the old ID table with the temporary ID table with compressed counts, step 144, before returning to step 120, FIG. 6.

Although this count compression routine somewhat reduces system performance, the

time to compute repeating data patterns and thus, to compress the "count" information in a record identification and locator table is minimal when compared to the tremendous savings of time which results from the ability to search the record locator table containing the count information in semiconductor memory instead of searching the disk drives for the requested record given the respective access times.

An example of a record identification and locator table for track 0 of a representative disk drive along with decoded information of each record is reproduced below as Table 4. This track identification and record locator table forms part of the device record identification and locator table as discussed previously in conjunction with FIG. 4. Line 2 of Table 4 corresponds to the track header portion 84, FIG. 5 of the track identification table also previously discussed in conjunction with table 3. The second byte of the header, the number "5" indicates that there are five records on this track.

TABLE 4

1.	TRACK NUMBER 0						
2.	FLAGS/COUNT/H/W/R/S/PAT/CACHE PTR: 00 5 00 00 00 00 0000000000000000						
00000000							
3.	FLAGS						
	00 01 03 03 01						
4.	BODY:						
	0000000000000000 0000000000000000 0000000000000000						
	0000000000000000						
5.	0000000000000000 0000000000000000 0000000000000000						
	0000000000000000						
6.	0000000000000000 0000000000000000 0000000000000000						
	0000000000000000						
7.	0000000000000000 0000000000000000 0000000000000000						
	000000000000000F						
8.	B0E050900418A0						
	RECORD FLAGS						
REC							
ID	KEY	DATA	WR	IN	NON PHYSICAL	ADDRESS	
#	CCHHR	LENGTH	LENGTH	PEND	CACHE	STD BLOCK	TRACK CACHE ADD
9.	00	0000000000	00	0008	...	00	00000000
							02E0
10.	01	0000000001	04	0018	...	01	00000001
							04E0
	02	0000000002	04	0090	...	03	00000002
							0800
	03	0000000003	04	0050	...	03	00000003
							0B80
	04	0000000004	00	0FB0	...	01	00000004

Line 3 begins the record flag portion of the track identification table and is comprised of five record flags namely, flags: 00; 01; 03; 03; and 01. Each of the record flags is associated with a corresponding record, in ascending order. Thus, record flag 00 is associated with data record 0; record flag 01 is associated with data record 1; and record flag 03 is associated with data record 2 and so forth.

A representation of a record modifier portion of the track identification and locator table is shown at lines 4-8 of Table 4. As discussed in conjunction with FIG. 5, the record modifier portion of the track identification and record locator table is read backwards beginning with the byte, "A0" of line 8.

The track identification and record locator table of Table 4 may be further understood in conjunction with lines 9-13. As shown on line 9, which identifies record 0 of this track, the second, third and fourth columns comprise the original "count" information of the CKD record. It should be noted that this record matches the description of the "expected" record utilized in the example associated with the method of FIG. 6 since the first record on the track is record 0, the key length is 0, and the data length is 8 bytes. Thus, the record locator flag associated with that record, "00" is the first record flag byte encountered on line 3.

Proceeding to line 10, record number 1 on the track has a key length of 04 and a data length of 18 and thus, deviate from the previously established "expected" data record and thus is assigned a record flag of 01.

Various codes which comprise the record flags are reproduced in Table 5 below wherein as shown in line 1, the code 00 means no change to the previously established "expected" record. As shown in line 2 of the table, record flag format 01 indicates that the first byte of the record modifier is the change flag byte, indicating that every bit flagged with a "1" points to the byte in the record identifier that should be replaced by the following bytes in the record modifier. The order of the record identifier is shown in line 7 of table 5 and begins with the key length, followed by data length (high), data length (low) and the first byte of the cylinder.

TABLE 5

RECORD FORM AT CHANGE CODES

1. ID FLAG FORMAT CODE 0:
NO CHANGE
2. ID FLAG FORMAT CODE 1:
1ST BYTE OF MODIFIER IS THE CHANGE FLAG BYTE
EVERY BIT FLAGGED POINTS TO BYTE IN THE
ID THAT SHOULD BE REPLACED BY THE NEXT INFO
BYTES.
- 3.
- 4.
- 5.
- 6.
- 7.
- 8.
- 9.
- 10.

NUMBER OF EXTRA INFO BYTES IS THE NUMBER OF '1's
IN THE 1ST BYTE. THIS CODE IS USED IF WE CAN'T
USE
ANY OTHER CODE.

:K.sub.L D.sub.LH D.sub.LL C CHRR

ID FLAG FORMAT CODE 2:
ONE BYTE INFO TO DL L + DL H = 2

ID FLAG FORMAT CODE 3:
ONE BYTE INFO TO DL L + DL H = 0

ID FLAG FORMAT CODE 4:
ONE BYTE INFO TO DL L DL H UNCHANGED

ID FLAG FORMAT CODE 5:
TWO BYTES INFO TO DL

ID FLAG FORMAT CODE 6:

ONE BYTE INFOR TO DL L + DL H 1

ID FLAG FORMAT CODE 7:
ONE BYTE INFO TO DL L + DL H = PATT FROM ID TABLE HEADER

ID FLAG FORMAT CODE 8:
DL L = PATT #0 FROM ID TABLE HEADER + DL H = 0

ID FLAG FORMAT CODE 9:
DL L = PATT #1 FROM ID TABLE HEADER + DL H = 0

ID FLAG FORMAT CODE A:
DL L = PATT #2 FROM ID TABLE HEADER + DL H = 0

ID FLAG FORMAT CODE B:
DL L = PATT #3 FROM ID TABLE HEADER + DL H = 0

ID FLAG FORMAT CODE C:
DL L = PATT #4 FROM ID TABLE HEADER + DL H = 0

ID FLAG FORMAT CODE D:
DL L = PATT #5 FROM ID TABLE HEADER + DL H = 0

20.
ID FLAG FORMAT CODE E:
DL L = PATT #6 FROM ID TABLE HEADER + DL H = 0

ID FLAG FORMAT CODE F:
DL L = PATT #7 FROM ID TABLE HEADER + DL H =

0

Thus, returning now to line 10 of table 4, the flag code 01 indicates that the first byte of the modifier namely, "A0" indicates the bits that are to be changed in the record identifier. Reading change byte A0 in conjunction with line 7 of Table 5 discloses that the successive bytes in the record modifier will modify the key length and data length (low) of the data record. The record modifier bytes in the track identification table modify the record identifier in reverse order as that shown in line 7 of Table 5 that is, from record number to key length. Thus, the second byte, "18" of the record modifier at line 8 of table 4 indicates that the previously expected data length is to be replaced with a data length (low) of "18", while the next byte of the record modifier, "04" is to replace the previously expected key length. It is in this manner that the system "reconstructs" the count portion of a CKD record from the "encoded" record identification and locator table.

Record number 2, line 11 of Table 4 also has a key length of "04" but the data length changes to "90". Thus, a flag of 03 is entered. The record flag of 03, as shown at line 9, Table 5, indicates that the next sequential byte of information in the record modifier is to be used as the data length (low) and the data length (high) will equal 0. Thus, the next consecutive entry of "90" in the record modifier portion of the track identification table body is accounted for.

Similarly, the next byte of the modifier portion of the track identification table is "50" which is the changed data length of record 3 read in conjunction with a record flag of "03" at line 12 of Table 4.

The final record flag "01" in the record flag portion of this track indicates that the next sequential byte namely, "E0" in the record modifier portion of the table is the changed flag byte pointing to the bytes in the record identifier that are to be changed or modified by the subsequent bytes in the record modifier portion of the table. Code E0 indicates that the key length, data length (high) and data length (low) are to be changed by the three bytes which follow as indicated by line 7, table 5. Thus, byte "B0" of the record modifier is used as the data length (low); byte "0F" is used as the data length (high) byte, and byte 00 modifies the former key length entry.

The building of a record identification and locator table in accordance with the present invention greatly reduces the amount of fixed block disk space required to store the "count" portion of a CKD formatted data record.

An additional example of a track level record identification and locator table is reproduced in table 6 below and is useful in showing an entry in the data pattern table previously described in conjunction with FIG. 7. Table 6 is a representation of a track identification table for an exemplary

TABLE 6

1. TRACK NUMBER D00
 2. FLAGS/COUNT/H/W/R/S/PAT/CACHE PTR: 11 81 00 00 00 C7 3E14181500FB9024
 03FDE000
 3. FLAGS
 00A320 28 202020XX XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX
 XXXXXXXXXXXXXXXXXX
 4. XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX
 XXXXXXXXXXXXXXXXXX
 5. XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX XX
 6. BODY:
 XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX
 XXXXXXXXXXXXXXXXXX
 7. XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX XXXXXXXXXXXXXXXXXX
 XXXXXXXXXXXXXXXX28

RECORD FLAGS

REC	ID	KEY	DATA	WR	IN	NON PHYSICAL	ADDRESS
#	CCHHR	LENGTH		LENGTH		PEND	CACHE
							STD BLOCK TRACK CACHE ADD

8.	00 00D0000000	00	0008	.	.	.	0 00047DB0 02E0
9.	01 00D0000001	00	0028	D.	D.	3	00047DB1 04E0 03FDE118
10.	02 00D0000002	00	0028	.	.	.	D. 0 00047DB2 0700 03FD3150
	03 00D0000003	00	003E	.	.	.	D. 8 00047DB3 0920 03FDE188
	04 00D0000004	00	003E	.	.	.	D. 0 00047DB4 0B60 03FDE1D0
	05 00D0000005	00	003E	.	.	.	D. 0 00047DB5 0DA0 03FDE218
	06 00D0000006	00	003E	.	.	.	D. 0 00047DB6 0FE0 03FDE260

On line 2 of table 6, the track "header" information is presented including the first byte "11" indicating that on this track, there is at least one record which is in cache, and at least one record which has a write pending, as previously explained in conjunction with Table 3. The second byte of the track header, "81" indicates there are 81 records in this track, while byte 5, "C7" is the CRC byte for this track. The next byte of the header, "3E" is the first byte of the data pattern table which extends for 8 bytes ending with "24". In this example, bytes 1-7 of the pattern table are not used, but are merely shown for illustrative purposes only. The

last four bytes of the ck header, "03 FD E0 00" is the cache beginning memory address at which any records from this track which are stored in cache are located.

Of particular interest in Table 6 is record 03 located at line 11. Since the data length, "3E" of record 3 is a deviation from the previously established data length "28", a record flag of other than 00 is expected, and thus the record flag "08" is entered. As can be compared from line 14 of the record flag codes in Table 5, record flag code "08" indicates that the data length (low) of this record identifier is to be loaded with pattern 0, the first pattern from the identification table header and thus, the "3E" pattern from line 2 of Table 6 is used as the data length for record number 3 when the system reconstructs the data record. The record flag, "28" which is shown and underlined on line 3 of table 6 also indicates that the data of this record is stored in cache. The cache address ascertained by adding up the cache memory starting address (line 2) contained in the header of the track identification table along with the length of any intervening data or key information stored in cache.

Modifications and substitutions of the present invention by one of ordinary skill in the art are considered to be within the scope of the present invention, which is not to be limited except by the claims which follow.

CLAIMS:

We claim:

1. Apparatus for retrieving at least one requested data record stored on a storage medium comprising:

means for receiving a plurality of data records, each of said data records including at least a record identification portion and a data portion;

means, responsive to said means for receiving, for transferring and storing said plurality of data records on at least one of said storage mediums;

means, responsive to said means for transferring and storing, for generating a corresponding plurality of recording locator indices, each of said corresponding plurality of record locator indices associated with one of said plurality of data records and a corresponding recording identification portion, for uniquely identifying the location of each of said plurality of data records stored on said at least one of said storage mediums;

memory storage means, for storing at least said plurality of record locator indices and the associated record identification portions in one record locator table;

means for requesting access to one or more of said plurality of data records, and for providing the data identification portion corresponding to said at least one requested data record;

means, responsive to said means for requesting access, for searching said record locator table stored in said memory storage means, for locating said data record identification portion and associated one or more record locator indices corresponding to said at least one requested data record; and

record retrieval means, responsive to said means for searching, for directly retrieving from said at least one storage medium said at least one requested data record as directed by said associated record locator indices.

2. The system of claim 1 wherein said record retrieval means and said means for transferring and storing retrieves and stores at least one data record in said memory storage means; and

wherein at least one of said record locator indices further includes means for indicating that the at least one data record associated with said at least one of said record locator indices is stored in said memory storage means.

3. The system of claim 2 wherein each of said record locator indices further includes a pointer to the memory address of the data record stored in said memory storage means;

4. The system of claim 2 wherein each of said record locator indices includes means for indicating that the associated data record is stored in said memory storage

means, and is awaiting waiting to one or more of said storage mediums.

5. The system of claim 1 wherein each of said record locator indices and associated record identification portions are combined to form said one record locator table stored in said memory storage means.

6. The system of claim 5 wherein said means for generating said record locator indices includes means for reducing the length of each of said record identification portions prior to storing said record identification portion and said corresponding record locator indices in said record locator table.

7. The system of claim 5 wherein said means for generating includes means for reducing the length of both said record locator indices and record identification portions prior to storing said record identification portion and record locator indices in said record locator table.

8. The system of claim 7 wherein said means for reducing transforms and encodes each of said record locator indices and record identification portions, for reducing the length of both said record locator indices and record identification portions prior to storing in said record locator table.

9. The system of claim 1 wherein said record retrieval means directly retrieves the one or more requested data records from said one or more storage mediums with no intervening search of said one or more storage mediums.

10. The system of claim 1 wherein said FBA formatted storage medium includes a directly addressable FBA formatted storage medium comprising directly addressable record storage locations.

11. The system of claim 1 wherein said FBA formatted storage medium includes a fixed block architecture (FBA) formatted disk drive.

12. The system of claim 10 wherein said directly addressable storage medium includes semiconductor memory.

13. The system of claim 10 wherein said directly addressable storage medium includes one or more optical disks.

14. The system of claim 10 wherein said directly addressable storage medium includes a CD ROM disk.

15. Apparatus for transforming and mapping CKD formatted data records onto fixed block disk drives, and for retrieving one or more requested data records stored on a fixed block disk drive utilizing a data record locator index stored in memory, comprising:

means for receiving a plurality of CKD formatted data records, each of said data records including at least a record identification portion and a data portion;

means, responsive to said means for receiving, for transferring and storing said plurality of data records on one or more fixed block disk drives;

means, responsive to said means for transferring and storing, for generating a plurality of record locator indices, each of said plurality of record locator indices associated with one of said plurality of data records, for uniquely identifying the location of each of said plurality of data records stored on said one or more fixed block disk drives, also including means for transforming and encoding each of said plurality of record locator indices and record identification portions, for reducing the length of each of said record identification portions and record locator indices;

means, responsive to said means for generating, for combining said transformed and encoded record locator indices and record identification portions, for forming a record locator table;

memory storage means, for storing said combined record locator table;

means for requesting access to one or more of said plurality of CKD formatted data records, and for providing one or more data identification portions corresponding to said one or more requested data records;

means, responsive to said means for requesting access to one or more of said

plurality of data records for searching said record locator table stored in said memory storage means, for locating said one or more data record identification portions and associated record locator indices corresponding to said one or more requested data records; and

record retrieval means, responsive to said means for searching, for directly retrieving from said one or more fixed block disk drives, said one or more requested data records as directed by said one or more record locator indices.

16. A method for transforming and encoding CKD formatted data onto one or more fixed block architecture (FBA) disk drives, and for retrieving one or more requested data records stored on one or more of said fixed block architecture disk drives utilizing a data record locator index stored in memory, comprising the steps of:

receiving a plurality of CKD formatted data records, each of said CKD formatted data records including at least a record identification portion and a data portion;

storing to one or more directly addressable FBA formatted storage mediums, said plurality of data records, said step of storing including:

transforming and encoding at least said record identification portion of each of said plurality of data records from a CKD format to an FBA format;

generating a plurality of record locator indices; and

combining said transformed and encoded record locator indices and record identification portions, for forming a record locator table, each of said plurality of record locator indices corresponding to one of said plurality of data records, for uniquely identifying the location of each of said plurality of data records stored on said one or more FBA formatted directly addressable storage mediums;

storing said combined record locator table in memory;

searching said record locator table stored in memory, for locating said one or more data record identification portions and associated record locator indices corresponding to one or more requested data records; and

directly retrieving from said one or more directly addressable storage mediums, said one or more requested data records as directed by said one or more record locator indices.

17. Apparatus for receiving one or more requested data records storage on a storage medium utilizing a data record locator index stored in memory, comprising:

means for receiving a plurality of data records, each of said data records including at least a record identification portion and a data portion;

means, responsive to said means for receiving, for transferring and storing said plurality of data records on at least one of said storage mediums;

means, responsive to said means for transferring and storing, for generating a plurality of record locator indices, each of said plurality of record locator indices associated with one of said plurality of data records and a corresponding record identification portion, and for uniquely identifying the location of each of said plurality of data records stored on said at least one of said storage mediums, said means for generating including means for transforming and encoding each of said record locator indices and record identification portions, for reducing the length of both said record locator indices and said record identification portions, and for combining said reduced length record identification portions with said corresponding reduced length record locator indices;

memory storage means, for storing at least said plurality of reduced length and combined record locator indices and corresponding record identification portions to form one record locator table;

means for requesting access to one or more of said plurality of data records, and for providing one or more data identification portions corresponding to said one or more requested data records;

means, responsive to said means for requesting access, for searching said record locator table stored in said memory storage means, for locating said one or more data record identification portions and associated record locator indices

corresponding to said one or more requested data records; and

record retrieval means, responsive to said means for searching, for directly retrieving from one or more of said storage mediums said one or more requested data records as directed by said one or more record locator indices.

18. The system of claim 17 wherein said plurality of data records and record identification portions are received in CKD format, and said plurality of data records are stored on at least said storage medium having an FBA format; and

wherein said means for reducing transforms and encodes said record identification portions from said CKD format to said FBA format.

19. Apparatus for retrieving at least one requested data record stored on a data storage medium, comprising:

a data record receiver, for receiving a plurality of data records, each of said data records including at least a record identification portion and a data portion;

a data record storер, responsive to said data record receiver, for transferring and storing said plurality of data records on at least one of said data storage medium;

a record locator index generator, responsive to said data record storер, for generating a corresponding plurality of record locator indices, each of said corresponding plurality of record locator indices associated with one of said plurality of data records and the corresponding record identification portion, for uniquely identifying the location of each of said plurality of data records stored on said at least one of said storage mediums;

memory storage, for storing at least said plurality of record locator indices and the associated record identification portions in one record locator table;

a data record access requestor, for requesting access to at least one of said plurality of data records, and for providing the data identification portion corresponding to said at least one requested data record;

a record locator table searcher, responsive to said data record access requestor, for searching said record locator table stored in said memory storage, for locating the data record identification portion and associated record locator indices corresponding to said at least one requested data record; and

a recording retriever, responsive to said record locator table searcher, for directly retrieving from said at least one storage medium, said at least one requested data record as directed by said associated at least one record locator indices.

WEST

L10: Entry 11 of 14

File: USPT

Dec 1, 1998

US-PAT-NO: 5845077

DOCUMENT-IDENTIFIER: US 5845077 A

TITLE: Method and system for identifying and obtaining computer software from a remote computer

DATE-ISSUED: December 1, 1998

INVENTOR- INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Fawcett, Philip E.	Duvall	WA		

ASSIGNEE- INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Microsoft Corporation	Redmond	WA			02

APPL-NO: 8/ 562929 [PALM]

DATE FILED: November 27, 1995

INT-CL: [6] G06 F 7/00

US-CL-ISSUED: 395/200.51; 364/221, 364/221.7, 364/222.81, 364/259, 395/712, 395/200.09

US-CL-CURRENT: 709/221; 717/173

FIELD-OF-SEARCH: 395/712, 395/200.09, 395/200.51, 395/232, 364/479.01, 380/4

PRIOR-ART-DISCLOSED:

U. S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>4796181</u>	January 1989	Wiedemer	364/406
<u>5047928</u>	September 1991	Wiedemer	364/406
<u>5142680</u>	August 1992	Ottman et al.	395/712
<u>5155484</u>	October 1992	Chambers, IV	
<u>5155680</u>	October 1992	Wiedemer	364/406
<u>5155847</u>	October 1992	Kirouac et al.	395/200.09
<u>5267171</u>	November 1993	Suzuki et al.	364/479.04
<u>5337360</u>	August 1994	Fischer	
<u>5367686</u>	November 1994	Fisher et al.	
<u>5388211</u>	February 1995	Hornbuckle	395/712
<u>5390247</u>	February 1995	Fischer	
<u>5421009</u>	May 1995	Platt	395/712
<u>5473772</u>	December 1995	Halliwell et al.	
<u>5495411</u>	February 1996	Ananda	395/232
<u>5528490</u>	June 1996	Hill	
<u>5548645</u>	August 1996	Ananda	380/4
<u>5586304</u>	December 1996	Stupeck, Jr. et al.	
<u>5586322</u>	December 1996	Beck et al.	707/200
<u>5654901</u>	August 1997	Boman	

OTHER PUBLICATIONS

Mori et al., "Superdistribution: The Concept and the Architecture", The Transaction of the Ieice, vol. E73, No. 7, pp. 1133-1146 (Jul. 1990).

Williams, "Internet Component Download," Microsoft Interactive Development, 49-52, Summer, 1996.

Rozenbilt, "O,A & M Capabilities for Switching Software Management" IEEE Global Telecommunications Conference, 1993, pp. 357-361.

ART-UNIT: 272

PRIMARY-EXAMINER: Swann; Tod R.

ASSISTANT-EXAMINER: Tzeng; Fred F.

ATTY-AGENT-FIRM: Klarquist Sparkman Campbell Leigh & Winston, LLP

ABSTRACT:

Creators of computer software provide the most up-to-date versions of their computer software on an update service. A user who has purchased computer software calls the update service on a periodic basis. The update service automatically inventories the user computer to determine what computer software may be out-of-date, and/or need maintenance updates. If so desired by the user, the update service computer automatically downloads and installs computer software to the user computer. By making periodic calls to the update service, the user always has the most up-to-date computer software immediately available. The update service may also alert the user to new products (i.e. including new help files, etc.), and new and enhanced versions of existing products which can be purchased electronically by a user from the update service.

24 Claims, 6 Drawing figures

WEST

End of Result Set

L6: Entry 4 of 4

File: USPT

Apr 21, 1998

US-PAT-NO: 5742792
 DOCUMENT-IDENTIFIER: US 5742792 A

TITLE: Remote data mirroring

DATE-ISSUED: April 21, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Yanai; Moshe	Brookline	MA		
Vishlitzky; Natan	Brookline	MA		
Alterescu; Bruno	Newton	MA		
Castel; Daniel D. C.	Framingham	MA		
Shklarsky; Gadi G.	Brookline	MA		
Ofek; Yuval Y. O.	Hopkinton	MA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
EMC Corporation	Hopkinton	MA			02

APPL-NO: 8/ 654511 [PALM]
 DATE FILED: May 28, 1996

PARENT-CASE:

RELATED APPLICATION This application is a continuation-in-part of U.S. patent application Ser. No. 08/052,039 filed Apr. 23, 1993, entitled REMOTE DATA MIRRORING (U.S. Pat. No. 5,544,347 issued Aug. 6, 1996), which is fully incorporated herein by reference.

INT-CL: [6] G06 F 12/16

US-CL-ISSUED: 395/489; 395/182.03, 395/182.04, 395/488, 395/492, 395/821, 371/10.1, 371/21.1, 371/40.1, 364/DIG.1

US-CL-CURRENT: 711/162; 710/1, 711/161, 711/165, 714/5, 714/6, 714/710, 714/718, 714/763

FIELD-OF-SEARCH: 395/182.03, 395/182.04, 395/488, 395/489, 395/492, 395/821, 371/10.1, 371/21.1, 371/40.1

PRIOR-ART-DISCLOSED:

U. S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>3588839</u>	June 1971	Belady et al.	395/469
<input type="checkbox"/> <u>3835260</u>	September 1974	Prescher et al.	379/237
<input type="checkbox"/> <u>3866182</u>	February 1975	Yamada et al.	340/172.5

Record Display Form

wysiwyg://59/http://westbrs:8002/bin/ga..._2=&p_doc_3=&p_doc_4=&p_doc_5=&p_doc_6=

<input type="checkbox"/>	<u>4020466</u>	April 1977	Cordi et al.	395/600
<input type="checkbox"/>	<u>4057849</u>	November 1977	Ying et al.	345/193
<input type="checkbox"/>	<u>4084231</u>	April 1978	Capozzi et al.	395/444
<input type="checkbox"/>	<u>4094000</u>	June 1978	Brudevold	345/200
<input type="checkbox"/>	<u>4124843</u>	November 1978	Bramson et al.	340/337
<input type="checkbox"/>	<u>4150429</u>	April 1979	Ying	395/293
<input type="checkbox"/>	<u>4161777</u>	July 1979	Behnase et al.	395/481
<input type="checkbox"/>	<u>4204251</u>	May 1980	Brudevold	395/309
<input type="checkbox"/>	<u>4342079</u>	July 1982	Stewart et al.	395/405
<input type="checkbox"/>	<u>4396984</u>	August 1983	Videki, II	395/858
<input type="checkbox"/>	<u>4430727</u>	February 1984	Moore et al.	395/442
<input type="checkbox"/>	<u>4453215</u>	June 1984	Reid	395/182.09
<input type="checkbox"/>	<u>4464713</u>	August 1984	Benhase et al.	395/415
<input type="checkbox"/>	<u>4577272</u>	March 1986	Ballew et al.	395/650
<input type="checkbox"/>	<u>4608688</u>	August 1986	Hanson et al.	371/11
<input type="checkbox"/>	<u>4634100</u>	January 1987	Julich et al.	371/11
<input type="checkbox"/>	<u>4698808</u>	October 1987	Ishii	371/21.5
<input type="checkbox"/>	<u>4710870</u>	December 1987	Blackwell et al.	395/182.04
<input type="checkbox"/>	<u>4755928</u>	July 1988	Johnson et al.	395/182.04
<input type="checkbox"/>	<u>4769764</u>	September 1988	Levanon	361/680
<input type="checkbox"/>	<u>4779189</u>	October 1988	Legvold et al.	395/493
<input type="checkbox"/>	<u>4783834</u>	November 1988	Anderson et al.	382/245
<input type="checkbox"/>	<u>4785472</u>	November 1988	Shapiro	379/96
<input type="checkbox"/>	<u>4797750</u>	January 1989	Karweit	358/335
<input type="checkbox"/>	<u>4805106</u>	February 1989	Pfeifer	395/650
<input type="checkbox"/>	<u>4814592</u>	March 1989	Bradt et al.	235/381
<input type="checkbox"/>	<u>4837680</u>	June 1989	Crockett et al.	395/284
<input type="checkbox"/>	<u>4849978</u>	July 1989	Dishon et al.	395/182.04
<input type="checkbox"/>	<u>4862411</u>	August 1989	Dishon et al.	395/494
<input type="checkbox"/>	<u>4916605</u>	April 1990	Beardsley et al.	395/489
<input type="checkbox"/>	<u>4949187</u>	August 1990	Cohen	358/335
<input type="checkbox"/>	<u>4985695</u>	January 1991	Wilkinson et al.	340/571
<input type="checkbox"/>	<u>5007053</u>	April 1991	Iyer et al.	371/21.1
<input type="checkbox"/>	<u>5029199</u>	July 1991	Jones et al.	379/89
<input type="checkbox"/>	<u>5051887</u>	September 1991	Berger et al.	395/489
<input type="checkbox"/>	<u>5060185</u>	October 1991	Naito et al.	395/600
<input type="checkbox"/>	<u>5089958</u>	February 1992	Horton et al.	395/182.03
<input type="checkbox"/>	<u>5097439</u>	March 1992	Partrquin et al.	395/402
<input type="checkbox"/>	<u>5099485</u>	March 1992	Bruckert et al.	395/182.09
<input type="checkbox"/>	<u>5123099</u>	June 1992	Shibata et al.	395/447

Record Display Form

<input type="checkbox"/>	<u>5127048</u>	June 1992
<input type="checkbox"/>	<u>5132787</u>	July 1992
<input type="checkbox"/>	<u>5134711</u>	July 1992
<input type="checkbox"/>	<u>5146576</u>	September 1992
<input type="checkbox"/>	<u>5146605</u>	September 1992
<input type="checkbox"/>	<u>5155814</u>	October 1992
<input type="checkbox"/>	<u>5155835</u>	October 1992
<input type="checkbox"/>	<u>5155845</u>	October 1992
<input type="checkbox"/>	<u>5157770</u>	October 1992
<input type="checkbox"/>	<u>5159671</u>	October 1992
<input type="checkbox"/>	<u>5170471</u>	December 1992
<input type="checkbox"/>	<u>5175837</u>	December 1992
<input type="checkbox"/>	<u>5175839</u>	December 1992
<input type="checkbox"/>	<u>5185864</u>	February 1993
<input type="checkbox"/>	<u>5201053</u>	April 1993
<input type="checkbox"/>	<u>5202887</u>	April 1993
<input type="checkbox"/>	<u>5206939</u>	April 1993
<input type="checkbox"/>	<u>5210865</u>	May 1993
<input type="checkbox"/>	<u>5235690</u>	August 1993
<input type="checkbox"/>	<u>5235692</u>	August 1993
<input type="checkbox"/>	<u>5239659</u>	August 1993
<input type="checkbox"/>	<u>5263154</u>	November 1993
<input type="checkbox"/>	<u>5269011</u>	December 1993
<input type="checkbox"/>	<u>5274645</u>	December 1993
<input type="checkbox"/>	<u>5276867</u>	January 1994
<input type="checkbox"/>	<u>5285451</u>	February 1994
<input type="checkbox"/>	<u>5313664</u>	May 1994
<input type="checkbox"/>	<u>5335352</u>	August 1994
<input type="checkbox"/>	<u>5343477</u>	August 1994
<input type="checkbox"/>	<u>5375232</u>	December 1994
<input type="checkbox"/>	<u>5377342</u>	December 1994
<input type="checkbox"/>	<u>5379412</u>	January 1995
<input type="checkbox"/>	<u>5381539</u>	January 1995
<input type="checkbox"/>	<u>5428796</u>	June 1995
<input type="checkbox"/>	<u>5446872</u>	August 1995
<input type="checkbox"/>	<u>5459857</u>	October 1995
<input type="checkbox"/>	<u>5463752</u>	October 1995
<input type="checkbox"/>	<u>5544345</u>	August 1996
<input type="checkbox"/>	<u>5544347</u>	August 1996

wysiwyg://59/http://westbrs:8002/bin/ga..._2=&p_doc_3=&p_doc_4=&p_doc_5=&p_doc_6=	
Press et al.	379/100
Omi et al.	358/524
Asthana et al.	395/800
Beardsley et al.	395/440
Beukema et al.	395/821
Beardsley et al.	395/872
Belsan	395/441
Beal et al.	395/182.04
Beardsley et al.	395/439
Iwami	395/250
Bonevento et al.	395/837
Arnold et al.	395/425
Ikeda et al.	395/410
Bonevento et al.	395/868
Benhase et al.	395/289
Ueno et al.	371/10.1
Yanai et al.	395/400
Davis et al.	395/575
Beardsley et al.	395/440
Ayres et al.	395/849
Rudeseal et al.	395/800
Eastridge et al.	395/575
Yanai et al.	395/425
Idelman et al.	395/182.04
Kenley et al.	395/600
Henson et al.	371/11.1
Sugiyama et al.	364/405
Yanai et al.	395/800
Yamada	395/182.02
Legvold et al.	395/575
Sakai et al.	395/425
Eastridge et al.	395/575
Yanai et al.	395/425
Iskiyan et al.	395/728
Ayres et al.	395/180
Ludlam et al.	395/182.04
Behnase et al.	395/481
Carpenter et al.	395/477
Yanai et al.	395/489

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	US-CL
0239323	September 1987	EPX	
0323123	December 1988	EPX	
57-111900	July 1982	JPX	
1-19437	January 1989	JPX	
1-19438	January 1989	JPX	
1-120650	May 1989	JPX	
2-32419	February 1990	JPX	
2-32420	February 1990	JPX	
2-32418	February 1990	JPX	
2-91716	March 1990	JPX	
2-91717	March 1990	JPX	
2-93721	April 1990	JPX	
2086625	May 1992	GBX	
PCT/US/93/05853	January 1994	WOX	
WO94/25919	November 1994	WOX	
PCT/US/84/01678	July 1995	WOX	

OTHER PUBLICATIONS

SFT Netware 286 Maintenance, #100-313-001, 100/Rev1.00, Novel Incorporated, Provo, Utah, Nov. 1987.

SFT Netware 286 Installation, #100-312-001, 100/Rev1.00, Novel Incorporated, Provo, Utah, Nov. 1987.

SFT Netware 286 Installation Supplement, #100-000225-001, 26/Rev1.02, Novel Incorporated, Provo, Utah, June 1986.

SFT Netware 68 Maintenance, #100-000289-001, 64/Rev1.00, Novel Incorporated, Provo, Utah, Nov. 1986.

SFT Netware 68 Installation, #100-000317-001, 103/Rev1.00, Novel Incorporated, Provo, Utah, Feb. 1988.

SFT Netware 868 Installation Supplement, #100-000286-001, 61/Rev1.00, Novel Incorporated, Provo, Utah, Dec. 1986.

D.L. Burkes & R.K. Treiber, "Design Approaches for Real-Time Transaction Processing Remote Site Recovery," Computer Society International Conference (COMPICON), Spring Meeting, Los Alamitos, Feb. 26-Mar. 2, 1990, No., CONF. 35, 23 Feb. 1990, Institute of Electrical and Electronics Engineers, New York, N.Y., pp. 568-572.

Supplementary European Search Report for EP 94 91 4223, Jul. 1, 1997.

International Search Report for PCT/US/94/04326, Jul. 20, 1994.

PCT Written Opinion for PCT/US94/04326, 28 Feb. 1995.

PCT International Preliminary Examination Report for PCT/US94/04326, Jul. 18, 1995.

Hank Cote-New Horizons in Accessing Large Amounts of On-Line Data, pp. 71-75, 1982.

Matt Kramer-Fault-Tolerant LANs Guard Against Malfunction, Data Loss, pp. C26-C35, Sep. 1987.

Carol Grossman-Planning for 3990 Extended Functions, pp. 245-258, 1989, 30th Annual GUIDE Conference.

Bob Buchanan-Administrative Error Causes Most Major System Failures (3 pgs.), Feb. 1988, Goverment Computer News, V. 17, N. 4, p. 340.

Virginia Dudek-Planning for Network Disasters (4 pgs.), Jun. 12, 1989, MIS Week, vol. 10, N. 24.

Jean Bozman-Escon An Aid To Remote Disks (3 pgs.), V. 25, N. 1, p. 29, Jan. 7, 1991.

Gary A. Boles-A RAID Stack: Micropolis RAIDION, Network Computer, Jun. 1992.

Blaine Homer-RAID Level 5 Protects Your Data With Interleaved Parity, LAN Times, May 25, 1992, vol. 9.

Michael Caton-Micropolis Array Secures Data, PC Week, 1992.

RAIDON Fault-Tolerant Disk Arrays Brochure, 1993.

StorageTek 20 Earnings Down, Iceberg Ready for Testing by Jim Mallory, Newsbytes, Jul. 15, 1993.

Network Backup Evolves, by Michael Peterson et al., PC Magazine, vol. 12, No. 16, Sep. 28, 1993, p. 277(18).

Your Next LAN Storage Could Be a Mainframe, Jeff Moad, Datamation, vol. 39, No. 6, Mar. 15, 1993, p. 71.

Heard Any Good Rumors Lately? (Network Management), By Jill Huntington-Lee, LAN Computing, vol. 5, No. 2, Feb. 1994, p. 22.

Special Report: The AT&T Worldwide Intelligent Network-~~Scope~~, Nature and Management, Edge, vol. 7, No. 224, Nov. 6, 1992, p. 9.

Protect Your Data! (Database Report), by Michael Liczbanski, Data Based Advisor, vol. 10, No. 5, May 1992, p. 114.

Network Planning to the Rescue; Network Managers Should Design and Test a Network Recovery Plan That Lets Business Continue Uninterrupted Even If Disaster Strikes, by Ramond W. Rudnitskas et al., Networking Management, vol. 9, No. 11, Oct. 1991, p. 46.

If LANS Eye Host Status, Quality Today Falls Short, by Barbara Bochenski, Software Magazine, vol. 11, No. 12, Oct. 1991, p. 66.

Safe and Secure, by Peter Stephenson, LAN Magazine, vol. 6, No. 9, Sep. 1991, p. 34.

Firing Up Fault Tolerance by Brian O'Connell, DEC Professional, vol. 10, No. 6, Jun. 1991, p. 118.

A Which Computer? Guide to Backup Systems, by Chris Bidmead, Which Computer?, Sep. 1989, p. 62.

Optical Disks Compete With Videotape and Magnetic Storage Media, by Henry Urrows et al., Optical Information Systems, vol. 8, No. 3, May-Jun. 1988, p. 101.

Preventing Disasters: Whipping Up an Emergency Response Plan, Network Computing, Feb. 1992, p. 81.

Abbey Banking on FDDI, Communications Week International, Mar. 4, 1991, p. 22.

X.25 Data Networks and IBM: A Troublesome Connection, by Jan Johnson, Computer Decisions, vol. 19, Jan. 12, 1987, p. 16.

DSC Communications Corp. Debuts NEXOS LAN Product Line at Interface, Demos Enhanced Performance, Fault Tolerance and WAN Connectivity, News Release, Mar. 28, 1988.

Analysis of Scanning Policies For Reducing Disk Seek Times, by E. G. Coffman, L. A. Klimko and Barbara Ryan-Siam, Journal on Computing, Sep. 1972, vol. 1, No. 3, pp. 269-280.

Configuring Tandem Disk Subsystems, by Scott Sitler, Tandem Systems Review, Dec. 1986, pp. 80-91.

Technical Report Investigation of Selected Disk Systems, Oct. 1976, Teledyne Brown Engineering, Huntsville, Alabama (18 pages).

Software Extends Unix Distributed Computing, by Evan O. Grossman, PC Week, Jun. 5, 1989, p. 35 and p. 40.

IBM Funds Pack Firm, Electronic News, May 22, 1989, p. 12.

IBM Plans U.S. Test of File-Transfer, Network World, Sep. 5, 1988, pp. 15, 17.

IBM Invests In Firm That Offers Unix File-Sharing, Network World, May 22, 1989, p. 4.

IBM Plans Speedy Nationwide File System, Computer Systems News, Sep. 5, 1988, p. 56.

Design Alternatives for Disk Duplexing, by Spencer W. Ng, IBM Research Report, Jan. 30, 1987.

Disk Arm Movement In Anticipation of Future Requests, by Richard P. King, IBM Research Report, Dec. 16, 1987.

Ubik: Replicated Servers Made Easy, by Michael Leon Kazar, pp. 60-67, 1989.

Carnegie-Mellon Forges Standard Look For Multivendor Net, Computerworld, Jan. 30, 1989, pp. SR/8-SR/9.

A Multi-Media Message System For Andrew, USENIX Winter Conference, Feb. 9-12, 1988, pp. 37-42.

Synchronization and Caching Issues in the Andrew File System, USENIX Winter Conference, Feb. 9-12, 1988, pp. 27-36.

Comparative Study of Some UNIX Distributed File Systems, EUUG, Sep. 22-24, 1986, pp. 73-82.

VOLUMES-The Andrew File System Data Structuring Primitive, EUUG Autumn 1986, Sep. 22-24, 1986, pp. 473-480.

Unix Plays Role in Corporate Networking, MIS Week, vol. 10, No. 28, Jul. 17, 1989, pp. 22-23.

Sun's Open System Made by Marketplace, "Unlike Minis, Unix Machines, PCs and Even VMS Can Be Glued" by NFS, Software Magazine, vol. 9, No. 12, Oct. 1989, pp. 72-75, Cashin, Jerry.

Universal Server Under Development, MacWEEK, vol. 3, No. 39, Oct. 31, 1989, p. 5.

Unix Plays Role In Corporate Networking-Vendors Look Unix-Based Network Management Tools as Networking Standards Slowly Evolve, MIS Week, Jul. 17, 1989, pp. 22-23, vol. 10, No. 28.

The 28 Runners Declared for Open Software Foundations, Distributed Computing Stakes, Computergram International, No. 1301, CGI11909912, Nov. 9, 1989.

Some Platforms (Theory of Self-Organizing Systems), Release 1.0, vol. 89, No. 6, Jun. 27, 1989, pp. 12-16.

Self-Organizing Systems (Overview of Detailed Examples of Self-Organizing Systems, Includes related article About the Importance of Self-Organizing Systems), Release 1.0, vol. 89, No. 6, Jun. 27, 1989, pp. 1-9.

IBM-Funded Startup Developing Unix Add-On, PC Week, vol. 6, No. 22, Jun. 5, 1989, p.

40.

Uniting File Systems; Experimental Large Scale, Distributed File Systems Are Now Being Created (included related article on the History of the AFS Project), *UNIX Review*, vol. 7, No. 3, Mar. 1989, pp. 61-70.

IBM Backs Transarc, Formed to Offer Andrew File System Distributed Unix Database, *Computergram International*, No. 1180, CGI05190001, May 19, 1989.

A Comparative Analysis of Disk Scheduling Policies, by Toby J. Reorey and Tad B. Pinerton, *Communications of the ACM*, Mar. 1972, vol. 15, No. 3, pp. 177-184.

Disk Scheduling: FCFS vs. SSTF Revisited, by Micha Hofri, *Communications of the ACM*, Nov. 1980, vol. 23, No. 11, pp. 645-653.

Using Multiple Replica Classes to Improve Performance in Distributed System, by Peter Triantafillou and David Tayloer, *The 11th International Conference on Distributed Computing Systems*, May 20-24, 1991, pp. 420-428.

Government Computer News, Oct. 9, 1987, Richard A. Damca, Microcomputing "Patch Fault Tolerance: A Strong Second Team," vol. 7, Issue 149 P27(2), Dec. 17, 1990.

Computer Reseller, Feb. 19, 1990, Joel Shore, Mirrored Servers' a reality, pp. 51, 52, 62.

Infoworld Articles, Mar. 12, 1990, Vendors Offer Mirroring Solutions P(31).

Wilkinson, Compact and Novell team for fault tolerance: mirrored servers to cut down on costly downtime, *Article MIS Week*, vol. 11, Issue 7, Feb. 12, 1990.

Article PC Week, vol. 4, Issue 37, Sep. 15, 1987, Matt Kramer, Fault-Tolerant LANS Guard Against Malfunction, Data Loss.

"The DASD Time Bomb," *Datamation*, vol. 32, Mar. 1, 1986, P63(4), Alper, Alan.

Symmetrix Remote Data Facility (SRDF) Product Guide, P/N 200-999-554 Rev A, EMC Corporation, Hopkinton, Mass., Sep. 1994, 105 pages [EMCP 00003686-379].

"Symmetrix Remote Data Facility Data Sheet, Symmetrix-resident Software Feature to extend mirroring capabilities," EMC Corporation, Hopkinton, Mass., Apr. 1995, 4 pages [EMCP 00010741-10741].

"When Data is Lost in Minutes, its an Act of God. When Data is Recovered in Minutes, its an Act of EMC." EMC Corporation, Hopkinton, Mass., 1994, one page [EMCP 00003903].

"EMC Ships Mainframe Disaster Recovery Solution that Restores Data in Minutes," News from EMC Corporation, EMC Corporation, Hopkinton, Mass., Oct. 3, 1994, two pages [EMCP 00003921-3922].

"Asset Protection: Disk-Based Approach Makes Disaster Recovery More Efficient, Ensures Business Continuity," EMC Corporate Information, EMC Corporation, Hopkinton, Mass., 5 pages [EMCP 00003916-3920].

"EMC Ships Continuously-Available Disk-Based Disaster Recovery Solution," News from EMC Corporation, EMC Corporation, Hopkinton, Mass., Oct. 3, 1994, 4 pages [EMCP 00003923- 8926].

"Symmetrix Remote Data Facility: Taking the disaster out of disaster recovery," Brochure, EMC Corporation, Hopkinton, Mass., Sep. 1994, 8 pages [EMCP 00003908-3915].

"Symmetrix Remote Data Facility: EMC redefines disaster recovery," Brochure, EMC Corporation, Hopkinton, Mass., Sep. 1994, 4 pages [EMCP 00003904-3907].

"EMC Introduces Symmetrix 5500 Continuous Operation Disk Storage for Mission Critical IBM Mainframe Market," News from EMC Corporation, Hopkinton, Mass., Nov. 10, 1992, 3 pages.

"MCI Deploys New Fast Data Transfer Technology from EMC," News from EMC Corporation, Hopkinton, Mass., May 22, 1995, 2 pages [EMCP 00007520-7521].

"EMC Agreement with Data Switch to Promote Long Distance Disaster Recovery," News from EMC Corporation, EMC Corporation, Hopkinton, Mass., Nov. 30, 1994, 2 pages [EMCP 00007516-7517].

"EMC and Comdisco to Benchmark Data Mirroring Capability," News from EMC Corporation, EMC Corporation, Hopkinton, Mass., Jan. 24, 1995, 2 pages [EMCP 00007518-7519].

"EMC Ships Mainframe Disaster Recovery Solution that Restores Data in Minutes," News from EMC Corporation, EMC Corporation, Hopkinton, Mass., Oct. 3, 1994, 2 pages [EMCP 00007514-7515].

N. S. Prasad, "IBM Mainframes: Architecture and Design," McGraw-Hill Book Company, New York, N.Y., 1989, pp. 58-73 and 182-196.

"The IBM 3390 Storage Control Reference," No. GA32-0099-04, International Business Machines Corporation, Tucson, Arizona, 1991, pp. 1-304.

Patterson et al., "A Case for Redundant Arrays of Inexpensive Disks (RAID)," Report No. UCB/CSD 87/391, Computer Science Division (EECS), University of California, Berkeley, California, Dec. 1987, pp. 1-24.

Patterson et al., "Introduction to Redundant Arrays of Inexpensive Disks (RAID)," COMPCON 89 Proceedings, Feb. 27-Mar. 3, 1989, IEEE Computer Society, pp 112-117.

Ousterhout et al., "Beating the I/O Bottleneck: A Case for Log-Structured File Systems," *Operating Systems Review*, vol. 23, No. 1, ACM Press, Jan., 1989, pp. 11-28.

Douglis et al., "Log Structured File Systems," COMPCON 89 Proceedings, Feb. 27-Mar.

3, 1989, IEEE Computer Society, pp. 124-129.

Rosemblum et al., "The Design and Implementation of a Log-Structured File System,"
ACM Transactions on Computer Systems, vol. 1, Feb. 1992, pp. 26-52.

ART-UNIT: 238

PRIMARY-EXAMINER: Asta; Frank J.

ATTY-AGENT-FIRM: Arnold, White & Durkee

ABSTRACT:

Two data storage systems are interconnected by a data link for remote mirroring of data. Each volume of data is configured as local, primary in a remotely mirrored volume pair, or secondary in a remotely mirrored volume pair. Normally, a host computer directly accesses either a local or a primary volume, and data written to a primary volume is automatically sent over the link to a corresponding secondary volume. Each remotely mirrored volume pair can operate in a selected synchronization mode including synchronous, semi-synchronous, adaptive copy - remote write pending, and adaptive copy - disk. Direct write access to a secondary volume is denied if a "sync required" attribute is set for the volume and the volume is not synchronized. If a "volume domino" mode is enabled for a remotely mirrored volume pair, access to a volume of the pair is denied when the other volume is inaccessible. In a "links domino" mode, access to all remotely mirrored volumes is denied when remote mirroring is disrupted by an all-links failure. The domino modes can be used to initiate application-based recovery, for example, recovering a secondary data file using a secondary log file. In an active migration mode, host processing of a primary volume is concurrent with migration to a secondary volume. In an overwrite cache mode, remote write-pending data in cache can be overwritten. Write data for an entire host channel command word chain is bundled in one link transmission.

38 Claims, 24 Drawing figures

WEST

L6: Entry 1 of 4

File: USPT

Jul 18, 2000

US-PAT-NO: 6092066

DOCUMENT-IDENTIFIER: US 6092066 A

TITLE: Method and apparatus for independent operation of a remote data facility

DATE-ISSUED: July 18, 2000

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Ofek, Yuval	Hopkinton	MA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
EMC Corporation	Hopkinton	MA			02

APPL-NO: 8/ 656035 [PALM]

DATE FILED: May 31, 1996

INT-CL: [7] G06 F 17/30

US-CL-ISSUED: 707/10; 707/204

US-CL-CURRENT: 707/10; 707/204

FIELD-OF-SEARCH: 707/204, 707/8, 707/10, 395/182.04, 395/182.13

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>4866611</u>	September 1989	Cree et al.	364/705.08
<input type="checkbox"/> <u>5093787</u>	March 1992	Simmons	705/33
<input type="checkbox"/> <u>5101492</u>	March 1992	Schultz et al.	395/575
<input type="checkbox"/> <u>5185884</u>	February 1993	Martin et al.	395/575
<input type="checkbox"/> <u>5206939</u>	April 1993	Yanai et al.	395/400
<input type="checkbox"/> <u>5235601</u>	August 1993	Stallmo et al.	371/40.1
<input type="checkbox"/> <u>5263154</u>	November 1993	Eastridge et al.	395/575
<input type="checkbox"/> <u>5317731</u>	May 1994	Dias et al.	395/600
<input type="checkbox"/> <u>5357509</u>	October 1994	Ohizumi	371/10.1
<input type="checkbox"/> <u>5392390</u>	February 1995	Crozier	345/335
<input type="checkbox"/> <u>5432922</u>	July 1995	Polyzois et al.	395/182.04
<input type="checkbox"/> <u>5434994</u>	July 1995	Shaheen et al.	707/201
<input type="checkbox"/> <u>5450577</u>	September 1995	Lai et al.	395/182.13
<input type="checkbox"/> <u>5469503</u>	November 1995	Butensky et al.	379/265
<input type="checkbox"/> <u>5495601</u>	February 1996	Narang et al.	395/600
<input type="checkbox"/> <u>5497483</u>	March 1996	Beardsley et al.	395/180
<input type="checkbox"/> <u>5537533</u>	July 1996	Staheli et al.	395/182.03
<input type="checkbox"/> <u>5544347</u>	August 1996	Yanai et al.	711/162
<input type="checkbox"/> <u>5579318</u>	November 1996	Reuss et al.	370/410
<input type="checkbox"/> <u>5596706</u>	January 1997	Shimazaki et al.	395/182.04
<input type="checkbox"/> <u>5615364</u>	March 1997	Marks	707/202
<input type="checkbox"/> <u>5625818</u>	April 1997	Zarmer et al.	707/104
<input type="checkbox"/> <u>5649089</u>	July 1997	Kilner	395/182.04
<input type="checkbox"/> <u>5673382</u>	September 1997	Cannon et al.	395/182.04
<input type="checkbox"/> <u>5740397</u>	April 1998	Levy	711/114

ART-UNIT: 277

PRIMARY-EXAMINER: Vonbuhr; Maria N.

ATTY-AGENT-FIRM: Herbster; George A. Pearson & Pearson

ABSTRACT:

A data network with a remote data facility for providing redundant data storage and for enabling concurrent access to the data for multiple purposes. A local data processing system with a data facility stores a data base and processes transactions or other priority applications. A second system, physically separated from the first system includes the remote data facility that normally mirrors the data in the first system. In an independent access operating mode, the second system is enabled to operate on the data within its data facility concurrently with, but independently of, the operation of the first system. On completion of the independent operation, the second system reconnects with and synchronizes with the first system thereby to reestablish the mirroring operation of the second system.

27 Claims, 8 Drawing figures